




# Artificial Intelligence Search



# Adversarial Search




# Single-Agent Trees



# Value of a State

Value of a state:  
The best achievable  
outcome (utility)  
from that state




Non-Terminal States:

$$V(s) = \max_{s' \in \text{children}(s)} V(s')$$

Terminal States:

$$V(s) = \text{known}$$

# Adversarial Game Trees




# Minimax Values

States Under Agent's Control:

$$V(s) = \max_{s' \in \text{successors}(s)} V(s')$$

States Under Opponent's Control:

$$V(s') = \min_{s \in \text{successors}(s')} V(s)$$



Terminal States:

$$V(s) = \text{known}$$

# Tic-Tac-Toe Game Tree



MAX (X)



MIN (O)





MAX (X)



MIN (O)


TERMINAL

Utility



# Adversarial Search (Minimax)

- ▶ Deterministic, zero-sum games:
  - ▶ Tic-tac-toe, chess, checkers
  - ▶ One player maximizes result
  - ▶ The other minimizes result
- ▶ Minimax search:
  - ▶ A state-space search tree
  - ▶ Players alternate turns
  - ▶ Compute each node's **minimax value**: the best achievable utility against a rational (optimal) adversary





# Minimax Implementation


def max-value(state):

    initialize  $v = -\infty$

    for each successor of state:

$v = \max(v, \text{min-value}(\text{successor}))$

    return  $v$



def min-value(state):

    initialize  $v = +\infty$

    for each successor of state:

$v = \min(v, \text{max-value}(\text{successor}))$

    return  $v$

$$V(s) = \max_{s' \in \text{successors}(s)} V(s')$$

$$V(s') = \min_{s \in \text{successors}(s')} V(s)$$

# Minimax Implementation (Dispatch)

```
def value(state):
```

if the state is a terminal state: return the state's utility

if the next agent is **MAX**: return **max-value(state)**

if the next agent is **MIN**: return **min-value(state)**


```
def max-value(state):
```

initialize  $v = -\infty$

for each successor of state:

$v = \max(v, \text{value}(\text{successor}))$

return  $v$



```
def min-value(state):
```


initialize  $v = +\infty$

for each successor of state:

$v = \min(v, \text{value}(\text{successor}))$


return  $v$

# Minimax Example






# Minimax Efficiency

- ▶ How efficient is minimax?
  - ▶ Just like (exhaustive) DFS
  - ▶ Time:  $O(b^m)$
  - ▶ Space:  $O(bm)$




# Minimax Properties




Optimal against a perfect player. Otherwise?

# Evaluation Functions



# Evaluation Functions


- ▶ Evaluation functions score non-terminals in depth-limited search



- ▶ Ideal function: returns the actual minimax value of the position
- ▶ In practice: typically weighted linear sum of features:


- ▶ 
$$Eval(s) = w_1 f_1(s) + w_2 f_2(s) + \dots + w_n f_n(s)$$

# Game Tree Pruning






# Minimax Example




# Minimax Pruning



# Alpha-Beta Pruning

- ▶ General configuration (MIN version)
  - ▶ We're computing the MIN-VALUE at some node  $n$
  - ▶ We're looping over  $n$ 's children
  - ▶  $n$ 's estimate of the childrens' min is dropping
  - ▶ Who cares about  $n$ 's value? MAX
  - ▶ Let  $a$  be the best value that MAX can get at any choice point along the current path from the root
  - ▶ If  $n$  becomes worse than  $a$ , MAX will avoid it, so we can stop considering  $n$ 's other children (it's already bad enough that it won't be played)
- ▶ MAX version is symmetric



# Alpha-Beta Implementation


$\alpha$ : MAX's best option on path to root  
 $\beta$ : MIN's best option on path to root

```
def max-value(state,  $\alpha$ ,  $\beta$ ):  
    initialize  $v = -\infty$   
    for each successor of state:  
         $v = \max(v, \text{value}(\text{successor}, \alpha, \beta))$   
        if  $v \geq \beta$  return  $v$   
         $\alpha = \max(\alpha, v)$   
    return  $v$ 
```

```
def min-value(state,  $\alpha$ ,  $\beta$ ):  
    initialize  $v = +\infty$   
    for each successor of state:  
         $v = \min(v, \text{value}(\text{successor}, \alpha, \beta))$   
        if  $v \leq \alpha$  return  $v$   
         $\beta = \min(\beta, v)$   
    return  $v$ 
```

# Alpha-Beta Pruning Properties

- ▶ This pruning has **no effect** on minimax value computed for the root!
- ▶ Values of intermediate nodes might be wrong
  - ▶ Important: children of the root may have the wrong value
  - ▶ So the most naïve version won't let you do action selection
- ▶ Good child ordering improves effectiveness of pruning
- ▶ With “perfect ordering”:
  - ▶ Time complexity drops to  $O(b^{m/2})$
  - ▶ Doubles solvable depth!
  - ▶ Full search of, e.g. chess, is still hopeless...
- ▶ This is a simple example of **metareasoning** (computing about what to compute)



*Thanks*